NEOTERIS

Securing Remote Access

WHITE PAPER

NOVEMBER 2002

Confidential - Do Not Distribute

This paper describes how the Instant Virtual Extranet (IVE) can increase overall network security. The first section describes the current security landscape and how security principles relate to remote access. The second section describes several potential points of attack that can be made more vulnerable by remote access solutions. The last section explains how using the Neoteris IVE to securely project select corporate resources and applications to remote users can reduce the risk of deploying remote access and increase overall network security.

Remote Access Security and the Neoteris IVE

Today's security landscape is vast and growing evermore complex. Even as security problems grow larger and more difficult to resolve, it's clear that the following security requirements have changed little since the advent of network computing:

- User identities must be authenticated.
- Resource requests must be authorized.
- Data must be protected from snooping or manipulation in transit.
- Transactions must be audited, logged, and protected against repudiation.
- Systems must be protected against intrusion.
- Systems must be available to users and reasonably easy to use.

What has changed is the way people use network resources. Today, users in all parts of an enterprise need to access network resources from remote locations. To support this type of usage, network administrators have created a "super-web" of interconnected systems. This super-web includes corporate and partner networks and remote access workstations, including corporate PCs and employees' own computers. Such a vast and organic web of systems exposes an enterprise because vulnerabilities in any node can compromise the entire system. Nevertheless, business needs have forced enterprises to open up their networks; delivering network resources and applications to mobile
workers, telecommuters, and remote partners is just too valuable to forego.

But the risk of attack remains and, as the super-system has grown, successful attacks have grown ever easier to generate. Malicious programmers have created tools that enable "script-kiddies" to initiate attacks. It's no longer feasible to tuck resources and applications safely inside a closed corporate network, and the risk and cost of deploying remote access is increasing, leaving enterprises with a dilemma:

1. **Accept with a higher degree of risk**
2. **Try to survive without ubiquitous connectivity**

By choosing one of the early approaches to providing this connectivity, approaches that focus purely on low-level network connectivity, enterprises expose resources to attack. By bringing the remote user inside the extended network with a VPN or by deploying applications in a segregated network partition with a custom extranet, enterprises introduce new costs and security risks. These costs and risks often offset the benefits of secure remote access, leaving network administrators unable to deploy access to their entire user community without incurring increased costs or exposing their network to unacceptable vulnerabilities.

The trend toward broad network access and the increase in losses due to security breaches have grown in parallel. The 2002 CSI/FBI Computer Crime and Security Survey clearly demonstrates how traditional remote access security has proven inadequate in protecting resources, applications, and systems. Even though most respondents use intrusion detection systems and firewalls, and even though almost all respondents use anti-virus software, the vast majority of respondents were still victimized by a security breach during the previous year. Attempts to close the security gap have fallen short. Token-based systems, PKI, and biometrics have improved authentication, but have complicated system integration and increased deployment costs. Complex and expensive access control systems have attempted to improve authorization, but still the attacks continue. They range from simple "cyber-vandalism" to more serious attacks that result in financial fraud or theft of proprietary information.

On average, during 2001, firms lost over six million dollars each due to security breaches, and a single firm was the victim of information theft that resulted in over 50 million dollars in damages. As ubiquitous network access has increased so has the cost of computer crimes and security breaches, with average annual losses increasing by over five million dollars since 1997. The total reported losses from inadequate protection of computers and network resources have totaled more than one billion dollars since 1997. Clearly, the network-centric approaches that either invite users into the LAN or place critical servers in public-facing network partitions need to be improved.

The facets of the remote access problem that these early implementations leave unmitigated must be addressed:

- VPN hardware might be relatively simple to install—it's an appliance deployed at the network perimeter. But, with a VPN, deployment is complex and costly. User connections must be provisioned and the users must be trained, so they understand the security risks inherent in having a long-lived network layer connection to the corporate network. If users cannot maintain the software configuration, the resulting downtime increases costs. Flaws in the client PC can also introduce security problems, especially if the client host is not secured with a personal firewall or with host integrity software.

- Custom extranet solutions, too, offer benefits—the usage model is simple and often no network-layer connection is required to access applications. But custom extranets are very expensive to deploy. While it may be inexpensive to add users to a Web application deployed in a DMZ, the initial cost of designing, developing, and deploying the extranet are very high. Furthermore, the hidden cost of maintaining a public-facing application server and the risk that flaws in the server software configuration will introduce security problems offset this ease-of-use benefit.

Neoteris developed the Instant Virtual Extranet appliance to protect enterprise network and information security while meeting the changing needs of today's networks. The IVE reduces the risk of enabling remote access, complementing existing network firewalls and intrusion detection systems, and completing the enterprise security landscape. Because it works seamlessly with existing enterprise security infrastructure and because it does not open persistent network-layer connections to potentially hostile client nodes, the IVE enables ubiquitous remote access without risking a compromise of network security.

Reducing risk is the key element to maintaining security. A solution that meets today's needs must simultaneously provide a broad range of access and protect against a broad range of attacks. The IVE achieves this design goal by delivering a secure appliance that supports application-layer remote access and protects enterprise applications, resources, and systems with the same security that exists on the enterprise network. Because the IVE leverages the Secure Sockets Layer (SSL) support in modern Web browsers and because it does not provision a persistent network-layer connection to internal resources, the IVE does not require installation and configuration of a client on remote PCs. The IVE solution focuses on providing additional layers of security and protecting the network against vulnerabilities that result from potentially hostile remote hosts and the proliferation of public-facing applications servers.

Authentication

Traditional extranet solutions often use static passwords stored in enterprise directories or application vendor-specific databases. While static passwords sometimes provide acceptable authentication for internal applications, they do not provide sufficient authentication for a remote access solution. Rewriting these internal applications or integrating with a strong authentication system is very costly. To meet the needs of security conscious deployments, the IVE supports the

RADIUS and LDAP protocols and X-509 client certificates. It can also provide strong authentication of user identities before allowing remote access to applications without requiring direct integration with the applications themselves. To ensure seamless and secure interoperability, the IVE integrates with authentication solutions and supports two-factor authentication or digital-certificate authentication from leading vendors, such as:

- ActivCard
- SecureComputing
- RSA Security
- VeriSign

The IVE establishes an SSL connection with a remote user through a Web browser and ensures that the channel is secure before any user credentials, such a two-factor passcodes or digital certificates, are transported over the internet. After a secure session is established, the IVE can securely accept additional application credentials and present them to a broad range of authentication systems, including:

- NTLM authentication
- HTTP authentication
- HTTP cookies

These credentials can optionally be stored on the IVE, so users do not have to re-enter them if they access the IVE from many different locations. If the system has been configured in this manner, a secure hash of the user credential is stored on the IVE's AES-secured file system. Further, the IVE can leverage a firm's investment in and seamlessly integrate with Windows or Netegrity-based single-sign on systems.

Authorization

Establishing and validating a user's identity is an important part of remote access security. After a user has established a secure session with the IVE, the system must also check whether that user has access to each resource or application requested. In many cases, enterprises need to provide different authorizations for remote access users. For example, some URLs or applications to which users have access on the LAN might not be appropriate for remote use. To meet these design requirements, the IVE supports a centralized authorization policy and access control lists to enable "remote access aware" augmentation of the access controls.

The IVE's authorization support can integrate with native permissions and with external authorization policies. For example, the IVE supports the native file authorization protocols for MS Windows and UNIX volumes. The IVE also preserves the native authorization for web applications and presents resources using the same access controls that exist to protect them on the LAN. Using the IVE's policies and access controls, an enterprise can enhance authorizations to web applications, file servers, client-server applications, terminal hosts, and messaging servers without installing any new software or reconfiguring existing servers. If an enterprise has invested in the Netegrity SiteMinder to protect Web-based resources, the IVE can integrate with that external system and act as a SiteMinder agent without any changes to the protected application or the Netegrity Policy Server.

Data Integrity

After verifying a user's claim of identity and checking that the user is authorized to access a resource or application, the IVE securely transports resources and application traffic across the Internet, to and from the origin server and the client host. Regardless of how the application or server is configured, the IVE secures outbound traffic by encapsulating it in SSL before it leaves the enterprise network and secures all potentially insecure inbound traffic before it leaves the client host. The IVE ensures that data transported between a remote client and the IVE is not monitored, recorded, or altered in transit.

The IVE can use SSL version 2 or 3, and can handle the most sophisticated cipher, 3-DES encryption, supported by modern Web browsers. While it is unlikely that an attacked would ever gain physical access to a deployed IVE appliance, the IVE also employs AES encryption at the system level to prevent the information from being compromised.

Auditing

In addition to support for SSL and digital certificates, which ensures that the endpoints are authenticated and protects data in transit, the IVE supports detailed, application-level logging. For example, administrators can:

- View the log on the appliance through the Administrator Console
- View SYSLOG messages on a external logging server
- Export the log for archiving and analysis on another system.

The IVE's logging support can help administrators identify attempted attacks, track transactions in Web applications, or, when used in concert with SSL and certificates, protect against repudiation of actions taken through the IVE. Because the IVE authenticates end points and protects data during transport, administrators can investigate activity that passes through the IVE by analyzing the logs.

Availability

To support availability beyond a single system, the IVE supports stateful system peering and native fail over. All system, group, and user data that is maintained by any IVE appliance in a multi-unit cluster is securely propagated to other IVE appliances in the system. If a system in a cluster fails, another system can take over and seamlessly handle susequent requests. This high-availability support enables administrators to deploy a secure system that does not have a single point of failure. In addition, high availability support allows administrators to protect their remote access solution against denial of service attacks and enables them to scale IVEs beyond single-appliance systems.

The Neoteris IVE also supports internal resiliency within each appliance. Each IVE system is monitored by an internal watchdog subcomponent that checks the health status of IVE processes as well as the status of external servers that it requires, such as authentication servers or LDAP directories. When the system encounters an error, such as a process that has stopped accepting messages or a server that has stopped responding, the watchdog attempts to remedy the problem and can send SNMP traps to network management systems so administrators can take action if the problem cannot be resolved automatically.

SSL vs IPSec

The IVE system, which uses SSL to authenticate endpoints and encrypt data, is often compared to Internet Protocol Security (IPSec) –based systems for remote connectivity. Sometimes such comparisons lead to a "Which protocol is more secure?" debate. In reality, these debates are not really relevant. These protocols achieve similar goals; they provide for secure key exchange and provide strong data protection during transport. Despite significant differences in the protocols, IPSec and SSL are actually quite similar at a high level. Both technologies effectively secure network traffic, and each has associated tradeoffs, which make them appropriate for different applications.

The IVE effectively mitigates the tradeoffs associated with SSL, and does not introduce the tradeoffs that spawn from IPSec. For example, SSL has been criticized because, in the past, each application had to be SSL-enabled. This tight integration often required development of new functionality and distribution of new software, which sometimes made SSL solutions less attractive than systems that worked by connecting networks at the operating system level. The IVE solves this problem by leveraging the SSL functionality that exists in Web browsers and enabling other applications to use it. With the IVE, applications can be secured on the fly, with no development effort and with no new software installation. Web applications, Java applets, client-server applications, and messaging clients can be SSL enabled automatically. The IVE accomplishes this support without making connections at a low level in the network stack and does not introduce the added security risk of such connections.

SSL has also been criticized because it enables ubiquitous access and because it is too easy to deploy to a broad range of end users. In practical terms, though, this is not a fair criticism. Ease of use encourages users to not circumvent security systems, and enables network administrators to secure more users on more systems without increasing costs. The poor usability and deploy-ability of traditional VPN systems is not a credible security benefit. The complicated client software required by them effectively limits the physical machine that connects to the network, but this is not a feature exclusive to traditional VPNs. The IVE can enforce the same limits through X.509 certificates, IP address filtering, and resource-level authorization policies. Further, the IVE's policies are easy to maintain and do not require client software, whereas a traditional VPN's complex, often poorly understood configuration procedures can lead to security flaws.

Though the protocol implementations differ greatly, and though the IVE's SSL approach does not require a software client, the two systems share many similarities. SSL and IPSec use the same technologies for user authentication; the leading solutions support RADIUS integration, two-factor systems, and X.509 certificates. Both IPSec and SSL can use two-factor authentication to authenticate users. Both IPSec and SSL can employ client-side X.509 certificates to effectively authenticate a user on a specific PC. SSL and IPSec can use the same 168-bit Triple DES encryption algorithm (3DES) and can optionally use others, such as 56-bit DES for IPSec and 128-bit RC4 for SSL, to achieve broader compatibility. While lesser SSL implementations can negotiate down to the least common denominator algorithm, the Neoteris IVE can require SSL version 3.0 and 128-bit encryption. Though they are defined outside of each security protocol, most SSL and IPSec solutions employ some form of access-control as well. Other arguments in the SSL vs IPSec debate include arguments over whether the protocols provide more or less flexibility and ubiquity—and whether the traits are good or bad. An enterprise should consider these arguments at an even higher level: providing remote access to a network can introduce unexpected costs and risk. The security solution that most effectively mitigates this potential risk should be the solution selected.

For example, IPSec solutions create a full network connection, which can result in security vulnerabilities. This risk can be mitigated, but only by maintaining host-integrity and security-posture software on the client machine and by protecting the network against a potentially hostile client that has the ability to open arbitrary network connections. In contrast, SSL solutions can reduce costs significantly because the deployment and maintenance of client software and the risk of hostile network clients both disappear.

More important than the question of which transport encryption protocol is "better" is the question: "Which security technology best fills the need for a remote access solution?" Since IPSec can be used to secure any IP traffic and SSL is limited to TCP traffic, IPSec is well suited for long-lived connections where broad and persistent, machine-to-machine network-layer connections are required, such as with site-to-site VPNs. SSL, on the other hand, is well suited for applications where the system needs to connect individuals to applications and resources. In the past, this requirement was often met through a persistent network-layer connection, but this connection was not the requirement. It was a side affect of an implementation that connected host machines and network gateways at the network level in order to connect users with the applications and resources that they need to use remotely.

These network-level connections have proven expensive to support and difficult to secure for end-user remote access solutions. A solution that can provide the benefits of a low-level connection, an "on the LAN" feel without rewriting or updating application software, with increased security and lower cost of ownership better meets the need of most end users. The IVE solution provides access to the applications that users need:

- Outlook Messaging
- Lotus Notes Messaging
- Client-Server applications, including terminal services
- Web applications and intranet content
- Java applet
- SSH/Telnet Hosts

And the IVE deploys with high system security, strong encryption, and integration with two-factor authentication, without introducing the usability challenges or management complexities of IPSec-based solutions.

Access

Authentication, authorization, auditing, data integrity, and availability—these are the pillars of security, but a remote access solution that supports each of these crucial components of security must also preserve simplicity and ease of use. Some remote access solutions sacrifice ease of use to preserve aspects of security. Others compromise security to achieve improved usability. The IVE is the only solution that simultaneously achieves the highest levels of security while maintaining ease of use for the end users and the administrators alike. With the IVE, enabling remote access is as simple as signing into a secure Web application.

This ease of use increases security. The IVE protects against flaws in end-user hosts and poorly configured home networks or partner networks, because it does not initiate a persistent network-layer connection. By accessing resources at the application layer, the IVE reduces the risk of providing access to these systems. Because the IVE greatly reduces an attacker's ability to hijack a user's connection and launch an attack on network services, it makes remote access to the following systems safer for the enterprise:

- Web applications
- CIFS and NFS file servers
- Client-server applications
- Lotus Notes servers
- Java applets
- SSH hosts
- Exchange servers

Remote Access Points of Attack

Traditional remote access solutions, custom extranets and traditional VPNs, are complex and often manually configured systems that leave networks and resources exposed to attack. Even though enterprises have deployed IDS, have manually hardened servers, and have attempted to protect systems with network layer firewalls, their resources are still vulnerable. These systems, the applications that run on them, and the information that the applications process are still at risk of compromise, forcing enterprises to accept a high level of risk when they deploy remote access. The IVE appliance, a closed, hardened system that connects users to resources and applications, reduces this risk without requiring server configuration or deployment of client-side security software.

Custom extranets that accept only web-based traffic and support only Web applications help solve the problem of non-secure clients, because they do not create a persistent network tunnel with client PCs. However, these deployments have limitations:

- Custom extranets are cost prohibitive.
- They support a limited number of applications.
- They increase the number of servers that must be deployed in a DMZ.

Custom extranets are also vulnerable to “cookie leaks,” when valuable information or meta-data used by an enterprise web application leaves the corporate network and is stored on every PC that each user uses to access the applications. The degree to which these “cookie leaks” pose a vulnerability varies widely across Web applications, but each Web application that is deployed in a DMZ might contain such leaks, or other security vulnerabilities, that increase the risk of deploying a custom extranet with each application that is added to the system.

VPNs, too, fail to protect against “cookie leaks” or other application specific vulnerabilities, and they can expose an even wider range of application vulnerabilities by effectively adding the client PC to the network. These LAN-connected remote PCs usually have no standard host-security software and are often used for other purposes. Those PCs that are outside the control of the enterprise IT staff, present an even greater opportunity to hackers. Home PCs, especially those with an “always-on” Internet connection and no, or a poorly configured firewall, provide an effective attack platform. They often lack anti-virus software or have long outdated antivirus data files. They also lack software that prevents the installation of or monitors for the presence of Trojan horse or “back-door” programs.

A non-technical PC user might download an apparently harmless program that lies in wait on the client PC and then attacks network resources during a VPN session. With more and more home users running server software, such as a personal Web server, a malicious program could spread in the same manner as Code Red, by compromising a default installation of such an application and attacking corporate servers during a subsequent VPN session.

The problem of potentially hostile clients connecting over a VPN is much broader than exposing corporate servers directly in a DMZ, but extranet servers are a growing point of concern. Extranet application servers have proliferated over the last several years and each one of these public-facing servers is a potential point of attack. Application vendors attempt to secure server software by offering dedicated proxies or “webified” front ends, and these efforts offer some protection, but they also complicate deployments and increase costs. As the DMZ grows, the firewall policies that separate it from the corporate LAN, the rule sets that define the network’s edge and the network’s core, have also grown complicated. A single vulnerability in a server’s OS, a Web server, or other application software, or a single error in a firewall policy can put the entire network at risk.

Because of the risk of network layer connections and because of the risk that viruses, worms, and Trojan horses pose even to vendor-specific gateways, enterprises need an additional layer of protection that partitions remote-access connections from the rest of the enterprise network. By partitioning the requests and by embargoing requests until they are authenticated, authorized, and logged and further ensuring that no content or meta-data leaks from the system and remains latent on the client PC, a dedicated application gateway, such as the Neoteris IVE, can protect against threats posed by potentially hostile remote clients.

Even with a mission-specific application-aware remote access gateway, some administrators worry about an attack that surreptitiously co-opts a user's session during the life of that session. The IVE system reduces the risk of this form of attack. Traditional VPNs must, in contrast, take measures to block simultaneous Internet and LAN connections to protect against split-tunneling attacks, because rogue software running on the PC can potentially exploit the network-layer connection to the LAN. The IVE prevents this sort of attack in many ways. Its HTTPS-based usage model reduces risk, because to anyone who could compromise a client machine this connection looks no different than a secure connection to order a book on Amazon.com, bid on an item on EBay, or do online banking.

This usage model forces an attacker to first take over the user's machine and then take over the internet browser before they could access confidential information or steal the session token. This sort of attack is very unlikely to succeed and would be obvious to the end user if it were successful. Further, the IVE employs additional protection against the Neoteris session to prevent unauthorized access to the system.

And finally, if the client system was compromised, as have many custom extranets and remote access VPNs, the IVE system does not allow arbitrary network-layer connections, as do VPNs, and it is not vulnerable to being "rooted" and used as an internal attack platform, as happens sometimes with servers.

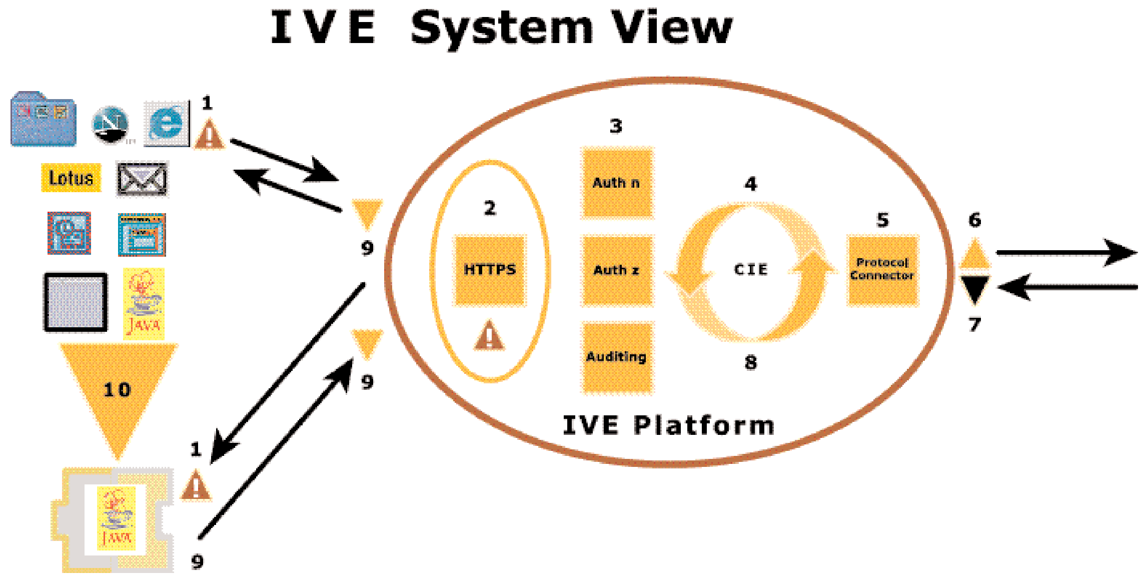

1. Request generated by the remote client

2. Request embargoed within the IVE Request Handler

3. Request validated within the IVE AAA subsystems

4. IVE Content Intermediation Engine transforms the inbound request

5,6,7. New request is generated and pass to the back end resource. Any response to the request is issued from the resource and sent to the IVE.

8. Outbound request is transformed by the IVE Content Intermediation Engine

9. Response sent to client

10. In the case of client/server applications, the client side Request Handler completes the transaction on the remote client.

Securing Remote Access with the IVE

To protect against the risks inherent by the very decision to deploy a remote access system:

- The IVE accepts only HTTPS connections, and no other inbound connections.
- Each connection is associated with a valid user credential, such as a two-factor pass code or a client digital certificate.
- All requests are authenticated based on the session token that is associated with this credential.
- When the browser connects to the IVE, the HTTPS request is embargoed within the IVE request handler.
- After checking authentication and authorization, the IVE connector makes a request on behalf of the end user.
- The headers and body of the response are returned to the request handler.
- The embargoed request terminates without ever reaching the internal network.

This process enables the IVE to handle remote requests without ever allowing the externally originated HTTPS request to reach internal server. The request handler is the only subsystem exposed to the requests, and it responds with the headers and body of the HTTPS response without ever accessing the internal resource itself.

As the single public facing component of the IVE system, the Request Handler deserves special scrutiny. The Request Handler is composed of two different public-facing subcomponents, the:

- Server-Side request handler: A specialized HTTPS daemon
- Client-Side request handler: A Java applet that secures application messages and sends them to the Server-Side Request Handler over HTTPS.

Server-Side Request Handler

The Server-Side Request Handler presents a hardened HTTPS server on a public facing network interface. The appliance resides behind a network firewall and terminates the HTTPS connections from web browsers and from instances of the IVE's Client-Side Request Handler. The IVE web server and request handler are the only components of the system that interact with the potentially hostile, externally originated HTTPS requests. The server has been hardened against URL crafting and parameter tampering attacks. Unauthenticated or unauthorized requests are logged and dropped before any attempt to access an internal server resource can begin. By embargoing the request and by initiating new requests on behalf of the user the server-side request handler protects against:

- Malicious, unauthenticated requests
- Authenticated, but unauthorized requests

In the unlikely event that an IVE session was compromised, for example if a PC with a valid session was left unattended in a public area, the server-side request handler still offers a measure of protection. In such a case it would still only allow access to those resources that the individual user was authorized to see and only for the life of the session. In contrast, an unattended VPN client could expose a broad range of servers to an attacker, including protocols and server resources that the individual user does not need or use.

Client-Side Request Handler

The IVE Secure Terminal, Client-Server and Messaging option use a subcomponent of the IVE request handler that is downloaded as a Java applet. This applet secures all IVE-bound network traffic before it leaves the remote host. The traffic is transported to the server-side request handler via HTTPS. This applet does not maintain direct network-layer connections with the server resources that the user is accessing through the IVE. There is no network adapter connecting directly to internal server addresses. Because the remote PC does not maintain a direct network connection, the network is not exposed to direct attack from hostile remote clients.

The IVE Client-Side Request Handler runs inside of the Java Virtual Machine (VM), where it is prevented from interacting with other programs and even other applets from other hosts. The "sandbox" of the Java VM ensures, even if the client host was hijacked, the applet cannot be hijacked by a Trojan horse or other malicious program unless that program can take over and "drive" the current instance of the browser or steal the IVE session token, which is protected by the cookie security model and by the Neoteris' dynamic transformation of JavaScript function bindings.

The Inadequacies of the "Fortress and Holes" Model

Direct LAN connectivity, resulting from remote-access VPN clients or the complex firewall forwarding tables required to deploy custom extranets weaken the network partitions that enterprises have put in place. Network administrators spend a great deal of time and money placing layers between the resources that reside at the center of the enterprise and the perimeter. The trend towards more ubiquitous remote access to mail and messaging servers, internal application servers, file servers, Citrix MetaFrame systems, and Microsoft Terminal Servers has led to increasingly common network compromises. By adding an application gateway to the network infrastructure that supports telecommuters, day-extenders, mobile workers, or business partners, an enterprise can continue to provide broad access without increasing risk. The Neoteris IVE adds a physical partition to the network as well as several logical layers of protection. It dramatically reduces the points of attack while simultaneously increasing an enterprise's ability to deploy a remote access solution to every employee or partner that has a business need for such access.

With the IVE, potential attackers see a single, hardened HTTPS server, and end users see a Web-browser window that automatically configures exactly, and only, the remote access that they need. In choosing the IVE system, network administrators reduce the need for high-cost, high-risk network-layer VPN connections, as well as custom extranet servers, which proliferate the risk in an already crowded and difficult to maintain DMZ

Instead they can deploy an HTTPS-based system that reduces the risk from client and server oriented attacks and makes it easy to support remote-access users. Such a solution must seamlessly integrate with an enterprises firewalls, IDS, and AAA infrastructure and mitigate the risks of remote access. A hardened appliance that connects users to the applications they need can secure remote access to enterprise networks and can help enterprises avoid network compromises that are growing more common and more costly. The IVE solution enables enterprises to deploy remote access that reduces costs and reduces risks.